Fermi National Accelerator Laboratory

FERMILAB-Conf-97/111

Commodity Hardware and Software Summary

Stephen Wolbers

Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

April 1997

Presented at *Computing in High Energy Physics*,
Lichtenberger Congress Center, Berlin, Germany, April 7-11, 1997

Operated by Universities Research Association Inc. under Contract No. DE-AC02-76CH03000 with the United States Department of Energy

Disclaimer

Distribution

Commodity Hardware and Software Summary

Stephen Wolbers

Fermi National Accelerator Laboratory, Batavia, IL

A review is given of the talks and papers presented in the Commodity Hardware and Software Session at the CHEP97 conference. An examination of the trends leading to the consideration of PC's for HEP is given, and a status of the work that is being done at various HEP labs and Universities is given.

Key words:

1 Introduction

There were 11 talks presented in the Commodity Hardware and Software Session, covering various issues of commodity computing. Seven of the papers concerned issues that affect the use of commodity computing for "physics" computing. This includes PC farms, compilers, linkers and multi-OS compatibility, and PC's for data acquisition. Four of the talks covered the issues of the administration of large numbers of PC's. This paper will summarize the talks that were given and will present some of my own views on some of the issues involved in commodity computing and its use in high energy physics.

Commodity computing is a little difficult to define but there are certain systems which clearly fall into this category. For example, a candidate list of commodity CPU processors would include the Intel Pentium and Pentium Pro, AMD, Cyrix, ALPHA, PowerPC and MIPS chips. This list is certainly not exhaustive. Based on volume the Intel Pentium and Pentium Pro are clearly commodity processors. There are many motherboards, graphics cards, Ethernet interfaces, etc. which can be purchased to build complete computing systems which include Pentium or Pentium Pro processors, and this allows one to build entire systems with commodity parts. In addition to the hardware there are many operating systems that can be considered candidate commodity operating systems. These include LINUX, SCO, FreeBSD, Solaris, Windows 95, Windows NT and OS/2. In the HEP community and in many others the choice has been narrowed down to LINUX or Windows NT, running on Pentium or Pentium Pro processors. The reason for this is clear. The

cost of these combinations for the computing power available is very low and the popularity of the systems is quite high, leading to excellent support. In the remainder of this paper commodity computers will be defined to be Intel Pentium or Pentium Pro based PC's, running either LINUX or Windows NT. (This includes any variants on the Pentium chip including MMX or Pentium II.)

2 Why Use Commodity Computing

The appetite for computing cycles in high energy physics (HEP) continues to rise rapidly, leading to a constant search for more cost-effective solutions for HEP computing needs. The trends in microprocessor technology lead us to notice that the Pentium Pro processor is quite a powerful chip, comparable in speed to many chips used in systems that HEP users are familiar with. The future expected improvement in speed will allow the Pentium chip (or its successors) to meet or exceed the absolute performance of the fastest workstation chips (ALPHA, MIPS, SPARC). This is very important, as there is a need for adequate single-processor speed in many applications. However, the real key to the attractiveness of PC's may be seen by examining the cost per unit of performance (the MIP) instead. PC systems tend to be very inexpensive when compared to workstations. Combined with the computing power of the PC's, the cost/MIP of computing on PC's is much less than workstations (or larger systems like SMP's) and the gap between them may grow.

There are real market forces which are pushing the PC's into the forefront. Millions of PC's are sold every year by many vendors, leading to economies of scale and severe price competition. This helps to drive the price down. The huge market also brings in large numbers of developers of software. This is truly a situation that the HEP community should not ignore.

3 Operating Systems

Once an Intel-based PC is purchased there remains a choice of which operating system to use. The two choices which are being considered by most people presenting papers in this conference are LINUX and Windows NT. These operating systems are both full 32 bit operating systems. The choice of which operating system to use depends on many factors. There are some real differences in the two operating systems and I will try to give a brief description of each.

First, I will give a short description of LINUX. LINUX is free. One can down-

load it off the network or if one desires can buy an inexpensive packaged version. The OS was originally written by Linus Torvalds, and is now supported by a large number of people coordinated by the GNU Project of the Free Software Foundation. Changes to the operating system occur quickly, with modifications based on bug reports, new devices or other new hardware, new features, etc. Many talks in this conference reported changes to LINUX which were made or expected within a day to correct problems that had been found in the course of working with the system. LINUX is a version of UNIX, meaning that all of the expertise built up in HEP over the years can be utilized and that all the features of UNIX can be exploited.

Windows NT is inexpensive. It is supported by Microsoft, meaning that there is a large organization devoted to developing and maintaining it. Windows NT is rapidly maturing, with release 4.0 appearing recently and version 5.0 expected in 1998. The newest versions are much more robust and feature-rich than previous versions. Windows NT is not UNIX, which to many people is an advantage. There are many software packages available for Windows NT, making the possibility of combining physics work with administration or other work on the same computers easily possible.

The ability to choose among these operating systems is good, as it gives flexibility to the HEP community. PC's can dual-boot, allowing one to change from LINUX to NT or vice versa. Both systems can be investigated and final decisions as to which to use can be deferred or possibly never made. Both operating systems have been shown to be effective for HEP and both will continue to be used in the near-future.

4 PC Administration

There were 4 talks given on the tools and techniques used to administer large numbers of PC's. This is a very large task, given that the HEP labs typically have 1000's of PC's to maintain. The cost of system administration is important and needs to be minimized. In addition there are certain services which people desire from PC administration, including a fairly uniform environment, good backup services, client-server implementation, license management and good software availability.

CERN continues to use and extend the NICE software systems with the recent addition of NICE95 and NICENT. These products have been quite successful in providing the kind of environment required. DESY is working on providing DESYNT, a client-server system which provides services to PC's in the DESY community. More and more PC's at DESY are taking advantage of this service to provide support. Fermilab is using Windows NT-Advanced Server to man-

age workgroups of PC's. Each workgroup has their needs defined by an MOU (memorandum of understanding). One goal is to have 75-100 (or more) PC's managed by one F.T.E. (full-time equivalent). The D0 model is a subset of the overall Fermilab model with the unique feature that it is primarily serving a group of physicists and is part of a collaboration-wide plan for computing.

All institutions are looking for ways to provide effective and less costly administration for PC's. This will continue as people look more closely at tools that Microsoft and others are providing and will provide in the future.

5 PC Farms

PC farms are an application of PC's that provide massive amounts of computing power. The PC farm is very similar to the standard UNIX workstation farm that HEP has been using for quite some time. The idea is to use dedicated sets of PC's as computing engines for CPU intensive tasks. The HERMES LINUX farm was described at this conference. This farm consists of 10 Pentium Pro 200 (soon to be two-processor machines (duals)) and has been in production reconstructing HERMES data. This farm is clearly a big success. The ZEUS NT farm consists of 18 Pentium Pro 200 PC's, running Windows NT. The ZEUS code has been ported and the system has been tested. Production should begin in spring or summer 1997. Fermilab has been investigating a PC LINUX farm, consisting of a mix of Pentium and Pentium Pro and dual PC's (9 PC's in all). Results of porting code and running parallel processing using CPS (cooperative processes software) have been encouraging. The CDF offline code has been ported and runs well on the PC farm. Finally, a group within the CLEO collaboration has been investigating a port of the NILE framework to PC's, running either NT or LINUX. A large production farm of approximately 50 nodes will be delivered this summer and real production should start to occur then.

One of the main conclusions to be drawn from the efforts that have occurred so far is that PC farms work, provide excellent performance, and provide it at an extremely low cost. It is likely that future acquisitions of computing meant to provide large CPU power will be some variant of PC farms.

6 Compilers, Utilities, Data Acquisition, etc.

To be able to effectively use PC's for HEP computing we must learn how to use the computers and in some cases how to maintain parallel code structures on both NT and UNIX. The CERN RD47 project and the NILE project have

both been investigating these issues. There are clearly compiler differences that must be understood and taken into account, for example. Both CERN and Fermilab have looked at a variety of compilers and at the differences in compile, link and execution speed coming from using these different compilers. The differences can be quite substantial.

A group at Rome has been working on PC's running LINUX for data acquisition. They have up to now been working on Camac connectivity to the PC and have had good success.

7 Conclusions

The conclusions drawn from the papers and talks given at CHEP97 on commodity hardware and software are the following. First, PC's are powerful and capable computers and can potentially supplant much of what is done on other machines today. Second, PC's for farms are clearly possible as was shown in the projects which are already underway and in production. Both NT and LINUX work well. Third, PC's on the desktop will likely supplant UNIX workstations. Finally it will be interesting to learn the progress that has been made in commodity computing in HEP at the next CHEP.